

Bears, all bears, and some bears.

Language Constraints on Language Models’ Inductive Inferences

Sriram Padmanabhan Siyuan Song Kanishka Misra
The University of Texas at Austin
{srirampadmanabhan, siyuansong, kmisra}@utexas.edu

Abstract

Language places subtle constraints on how we make inductive inferences. Developmental evidence by Gelman et al. (2002) has shown children (4 years and older) to differentiate among generic statements (“Bears are daxable”), universally quantified NPs (“all bears are daxable”) and indefinite plural NPs (“some bears are daxable”) in extending novel properties to a specific member (*all* > *generics* > *some*), suggesting that they represent these types of propositions differently. We test if these subtle differences arise in general purpose statistical learners like Vision Language Models, by replicating the original experiment. On tasking them through a series of precondition tests (robust identification of categories in images and sensitivities to *all* and *some*), followed by the original experiment, we find behavioral alignment between models and humans. Post-hoc analyses on their representations revealed that these differences are organized based on inductive constraints and not surface-form differences.

1 Introduction

A hallmark of human cognition is our ability to make inductive inferences about categories (Osherson et al., 1990; Murphy, 2004). We can readily extend immediately available information along different categorical structures—e.g., on learning that *robins* have a novel property, we might extend it to other *birds*. An active area of research in cognitive science is to understand exactly how these inferences are constrained (Gelman et al., 2002; Hollander et al., 2002; Cimpian and Markman, 2008). For instance, a goldfish is a pet, a fish, an animal, and a living thing, and so deciding the extent to which properties generalize beyond available evidence, to other categories, is a non-trivial task—one which we master anyway.

Language has been posited to play a particularly indispensable role in constraining category-based inductive inferences (Prasada, 2000; Gelman,

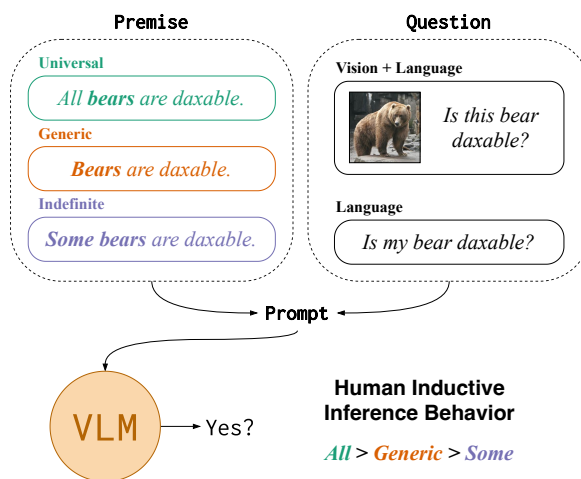


Figure 1: We study how the form of a proposition constrains a model’s inductive inferences. We compare between universal quantifiers (*all*), generics (*bare plurals*), and indefinite quantifiers (*some*). Humans represent them differently and show consistent graded effects in their inductive inferences (Gelman et al., 2002).

2004). One particular way in which language constrains inductive inferences is by modulating the scope of a proposition—e.g., (1a) expresses a completely different inductive constraint (i.e., the extent to which it will generalize) about the novel property “has the T9 hormone” than does (1b).

- (1) a. All bears have the T9 hormone. [ALL]
- b. Some bears have the T9 hormone. [SOME]
- c. Bears have the T9 hormone. [GENERIC]

While quantifiers like *all* and *some* are logically obvious in terms of the a priori expectations about how they might constrain a learner’s inductive inference, the status of statements like (1c) has been particularly puzzling from a theoretical standpoint. These statements—called *generics* (Leslie, 2008)—denote law-like *generalizations* about kinds, and are robust to counter-examples (encountering an unfriendly dog does little to alter one’s generic knowledge that “*dogs are friendly*”). Generics are

notoriously difficult to detect (Gelman, 2004), especially because unlike quantifiers like *all* and *some*, generics are never explicitly marked in any known language. Simply treating bare plurals as generics is insufficient (e.g., they can also be denoted using indefinite articles—*a goldfish has bad memory*) and sometimes incorrect (e.g., *Mosquitoes are torturing me*). Nevertheless, children can produce generic sentences before more theoretically and logically tractable ones involving quantifiers like *all* (Leslie, 2008). Experimentally, Gelman et al. (2002) showed children (4 years old) and adults to represent generics like (1c) differently than they did *all* and *some* in their inductive behavior. Participants maximally extended a property to a newly encountered bear when told that *all bears have a property*, followed by the bare plural *bears have a property*, followed by *some bears have a property*—i.e., their generalization behavior consistently showed the following pattern: ALL > GENERIC > SOME.

The puzzle of generics and its relation to quantifiers like *all* and *some* is quite relevant to modern artificial intelligence models like (vision) language models (VLMs). Apart from being central to the study of inductive reasoning in humans (a core property of our intelligence), the differences in the three kinds of propositions also allows us to shed light on how language *form* (all/some/generic) interplays with language *function* (inductive behavior) within these models. This is particularly relevant in light of recent positions about the ability of such models to grasp the nuance of meaning (Bender and Koller, 2020; Mahowald et al., 2024). Therefore in this paper, **we investigate the extent to which VLMs distinguish between ALL, SOME, and GENERICS in their inductive inferences**. We do so by replicating Gelman et al. (2002)’s experiment on VLMs.¹ We provide models with premises consisting of statements like those in (1) that attribute a property (*have the T9 hormone*) to categories (*bear*) and then test if they extend the property to a specific member of the category, presented to them as an image (like in the original experiment) or using other linguistic cues (e.g., *my bear*).

Before analyzing models’ inductive reasoning behavior, we first test if they satisfy basic presuppositions to our research question, by tasking them with the same pretests as in Gelman et al. (2002). This is especially important in order to

make “species-fair” comparisons between models and humans (Firestone, 2020). In particular, we test for two presuppositions. First, the models must recognize that their input image consists of the category that the premise and the subsequent question focuses on. For this, we test their ability to robustly identify categories in multiple object-centric images. Second, the models should capture the basic linguistic properties of *all* and *some*—both of which are theoretically and logically more tractable than generics. For this, we create a **new developmentally inspired benchmark** that tests if models can answer questions about a context presented to them in either modality (language vs. vision) targeting *all* and *some*—e.g., if the model says ‘yes’ to the question “Are *all* blocks on the table *blue*?”, pertaining to an image consisting of five blue blocks on a table and nothing else. Both these tests mimicked the pre-tests conducted by Gelman et al. (2002), and the *all* and *some* test in particular was also inspired by a classic experiment testing for quantifier knowledge in children of ages 3–4 (Smith, 1980). We ran these tests on 10 different VLMs, finding at least two that robustly satisfied both these presuppositions.

Turning to our inductive behavior results, we found both VLMs from the previous experiments to show **qualitatively similar patterns of inductive reasoning as humans**. That is, they were more likely to generalize a novel property to a category when the premise contained a universal quantifier *all*, followed by the bare plural *generic*, followed by the indefinite quantifier *some*. Our post-hoc analyses on models’ vector representations of only the premises, without any task context, revealed that these proposition types occupied different regions in the models’ low-dimensional representational space. This was found to hold even when we added stimuli that expressed the same proposition but differed in their surface forms to the original stimuli—e.g., adding *every bear* to pair with *all bears*. That is, models organized propositions **in terms of similarities in their inductive constraints, in a manner that cannot be explained by surface-form similarities/differences alone**.

Overall, our experiments and results contribute to the broader literature surrounding (V)LMs and their conceptual knowledge. Existing work on models’ inductive reasoning capacities has largely focused on the effect of the conceptual content of the premises (Misra et al., 2021, 2022; Han et al., 2024; Bhatia, 2023). For instance, Han et al. (2024) use

¹Our reason for using VLMs as opposed to LMs is because the original experiment by Gelman et al. (2002) was multimodal, and involved pictures of animals in its stimuli.

stimuli from Osherson et al. (1990), which does not contain any variation in terms of scope-modifying cues. We contribute to this body of work by narrowing our focus on the constraints imposed by various surface form changes, keeping the conceptual content the same (i.e., same concepts). Furthermore, our work also complements recent work on generics and (V)LMs. For instance, recent work has focused on the extent to which these models are able to reason *about* generics in the context of their interplay with exceptions (Allaway et al., 2023, 2024; Frank and Allaway, 2025). While generics have been explicitly compared against quantifiers (Ralethe and Buys, 2022; Collacciani et al., 2024; Cilleruelo et al., 2025), no work has investigated the differences between generics and quantifiers in terms of their inductive constraints, which is an important component of their meaning. Finally, an unintended contribution of our work is our developmentally inspired benchmark focusing on models’ behavior on *all* and *some*, which we release in both the language and the vision+language modalities, and is the first of its kind, to the best of our knowledge—with previous work instead focusing on vague quantifiers (Wong et al., 2025). This joins other developmentally inspired benchmarks for VLMs (Tan et al., 2024; Yiu et al., 2025). Our code is available at <https://github.com/kanishkamisra/inductive-constraints>.

2 Models and measures

Models Our experiments use stimuli that span both the vision and text modalities, and therefore we conduct our analyses on Vision Language Models. We test on 10 different models: Idefics3-8B (Laurençon et al., 2024), LLaVA-1.5 7B, LLaVA-v1.6-Mistral 7B (Liu et al., 2023b,a), LLaVA-OneVision (Li et al., 2025), Molmo 7B (Deitke et al., 2025), Qwen2.5-VL 7B Instruct (Qwen Team, 2025a) and Qwen3-VL Instruct 2/4/8B (Qwen Team, 2025b), and SmolVLM Instruct (Marafioti et al., 2025, 2.2B). Models were accessed using the `minicons` library (Misra, 2022). Appendix A shows details of the VLMs.

Measures Our stimuli in all three subsequent experiments involve a polar question, and expect models to answer with Yes or No. To account for surface form variation, we follow Rodriguez et al. (2025) and take into account space prefixing as well as both upper and lower cased versions of Yes and No into our mea-

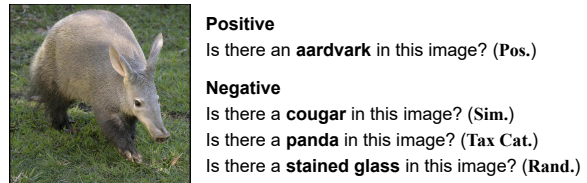


Figure 2: An instance from our category identification experiment. Each positive sample is associated with three separate types of negative samples.

asures. Given context C and a question Q , with $\text{YES} = \{\text{‘Yes’}, \text{‘yes’}, \text{‘ Yes’}, \text{‘ yes’}\}$, and $\text{NO} = \{\text{‘No’}, \text{‘no’}, \text{‘ No’}, \text{‘ no’}\}$ we compute the relative probability of Yes as:

$$p_{\text{rel}}(\text{Yes}) = \frac{\sum_{l \in \text{YES}} p_{\text{LM}}(l \mid Q, C)}{\sum_{l \in \text{YES} \cup \text{NO}} p_{\text{LM}}(l \mid Q, C)}, \quad (1)$$

3 Experiment 1: Category Identification

Our first experiment targets VLMs’ ability to robustly detect the presence (or absence) of categories in visual environments. This is an important precondition for our inductive generalization experiment, since a substantial part of that experiment involves testing models’ generalization of a novel property to a category indicated using an image. We measure a model’s robustness in category identification using two considerations: 1) by sampling multiple different ‘negative’ categories (ones that do not appear in the image) for each instance, ensuring that the model not only detects the presence of a given category, but also the absence of those not in the image; and 2) by measuring accuracy as a function of multiple images, and only ‘rewarding’ models iff. they are able to correctly perform category identification for *all* images.

3.1 Data

Our images are sourced from THINGS (Hebart et al., 2019), a resource of high quality, entity centric images collected for psychological experiments. THINGS consists of images across 1854 categories, with each category being paired with 10-20 images. We use a subset of 1222 categories from the database, and sample 5 images per category to measure image-level robustness. We use polar questions as our stimuli, with template “Is there a [CATEGORY] in this image? Answer with Yes or No.” Our positive samples (**Pos.**) are constructed by substituting [CATEGORY] with the appropriate lexical item as provided by THINGS. Following Misra et al. (2023) and Qin et al. (2025),

Model	Pos.	Sim.	Tax Cat.	Rand.	Joint
Llava-1.5-7B	97.5	26.6	52.2	78.0	15.0
Llava-OV-7B	35.8	58.9	79.0	90.0	17.4
Llava-1.6-7B	91.0	40.5	69.5	90.5	26.5
SmolVLM	80.1	58.0	82.4	95.0	37.8
Idefics3-8B	98.4	51.8	78.4	93.8	40.1
Molmo-7B	85.0	71.6	88.0	94.8	52.9
Qwen3-VL-2B	95.2	71.8	89.1	96.1	60.0
Qwen3-VL-4B	<u>98.3</u>	72.4	90.5	<u>97.3</u>	63.8
Qwen3-VL-8B	95.4	<u>75.4</u>	91.4	97.0	<u>65.2</u>
Qwen2.5-VL-7B	93.3	77.5	92.1	97.5	66.2

Table 1: Category Identification Results. Accuracies of VLMs on positive samples (**Pos.**), and the three negative sample subsets (**Sim.** = Similarity-based; **Tax. Cat.** = Sampling from same taxonomic category; **Rand.** = Random sampling). **Joint** indicates the proportion of time a model correctly answered all 20 questions for a given category. **Bold** and underline indicate best and second-best performance, respectively.

we pair each positive sample with multiple different negative samples, each coming from a different knowledge source (similarity, taxonomic relations, etc.). We consider three different sources:

SPoSE Similarity (Sim.) Our first class of negative samples uses SPoSE embeddings (Zheng et al., 2019; Hebart et al., 2023), a vector space model built using the THINGS images, which provides a high-fidelity estimate for visual similarity judgments from humans (Spearman’s correlation of 0.94, see Kaniuth et al., 2024). For each target concept, we perform weighted random sampling of negative concepts, with the weights proportional to the SPoSE similarity between concepts.

Sampling from taxonomic category (Tax Cat.) Our second class of negative samples comes from randomly sampling members of the smallest taxonomic category/hypernym that the target concept belongs to, as provided by Hebart et al. (2023) in the THINGS database. For instance, if *sparrow* is our target, then we sample from the *bird* category, expecting concepts like *crow*, *eagle*, etc. This strategy captures taxonomically related concepts.

Random sampling (Rand.) Finally our third class of negative samples comes from a simple random sampling over the other THINGS categories.

Since there are 5 unique images per positive concept, and 3 total negative samples, we end up with a total of 24,440 image, question pairs.

3.2 Results and Analysis

To recap our stimuli, each category is associated with 5 total images, and for each image, we have 4 types of questions—one positive question which asks for the presence of the category (expecting ‘Yes’), and three negatives, which ask for the presence of other category (expecting ‘No’). A competent model must be able to identify not only that the object in *all* images (per category) is the target category but also that it is not some other category. This results in two types of measures. First, for each question type, we measure the average percentage of cases the model correctly answers the question correctly for *all* images for a given category. We report this accuracy measure per question-type. Additionally, we also measure the average proportion of time the model correctly answers *all* four questions for *all* five images in the given category—i.e., all 20 questions associated with a given category. We refer to this latter measure as ‘**Joint**’ accuracy. Table 1 shows the accuracy measure across all question types as well as the Joint accuracy.

We see that VLMs achieved generally high accuracies on positive questions and randomly sampled negative questions. Their performance declined on negative questions constructed using concepts drawn from the same taxonomic category, though not by all that much. Finally, across all models, the lowest accuracy was observed on similarity-based negative samples. This insensitivity to similarity-based samples corroborates previous work on LMs’ ability to predict features (Misra et al., 2023). Out of all evaluated VLMs, Qwen2.5-VL-7B attained the highest **Joint** accuracy across all four sample types, with Qwen3-VL-8B being a close second.

4 Experiment 2: Behavioral Sensitivity to *All* and *Some*

Our next experiment focuses on models’ representation of the universal quantifier *all* and the indefinite quantifier *some*. We use this as a pretest to establish the minimal linguistic knowledge of the models to complete the task in our inductive inference experiments. This test specifically targets if models realize *all* as a quantifier to denote a scenario where every item had an attribute (e.g., color) and *some* as a quantifier where at least one item had the attribute. The test follows directly from Gelman et al. (2002), as well as classic experimental work on children’s ability to learn knowledge of quantifiers (Smith, 1980), who showed that children begin to learn the







Condition	Vision + Language	Language	Answer
All/All/All	 Are all cubes in the image blue in color?	There are 20 blocks on a table. 20 of them are red . Are all of the blocks on the table red ?	Yes
All/All/None	 Are all cubes in the image orange in color?	There are 20 blocks on a table. 20 of them are red . Are all of the blocks on the table blue ?	No
Some/All/Maj.	 Are all cubes in the image blue in color?	There are 20 blocks on a table. 12 of them are red and 8 of them are blue . Are all of the blocks on the table red ?	No
Some/All/Min.	 Are all cubes in the image orange in color?	There are 20 blocks on a table. 12 of them are red and 8 of them are blue . Are all of the blocks on the table blue ?	No
Some/Some/Maj.	 Are some cubes in the image blue in color?	There are 20 blocks on a table. 12 of them are red and 8 of them are blue . Are some of the blocks on the table red ?	Yes
Some/Some/Min.	 Are some cubes in the image orange in color?	There are 20 blocks on a table. 12 of them are red and 8 of them are blue . Are some of the blocks on the table blue ?	Yes

Table 2: Stimuli for testing knowledge of *all* and *some*, across input modalities: 1) **Vision + Language** ($N=840$), where questions target properties of entities in an image; and 2) **Language** ($N=900$), where questions target properties of entities in a math word-problem. There are a total of 6 conditions, varying in the properties of the input context (All/Some), the question (All/Some), and the amount of entities to which the property in the question targets (All/None/Majority/Minority). **Note:** stimuli in the two modalities are *not* matched and generated separately.

difference between *all* and *some* at age 3. However, since there currently is no such dataset that allows us to test this directly, we resort to creating our own tightly controlled benchmark consisting of stimuli in both vision and textual modalities.

4.1 Data

We design separate stimuli across visual and textual modalities. Nonetheless, our stimuli in both modalities follow a general template consisting of a context (visual or textual) that presents objects with one or more properties, followed by a question that inquires about the properties of these objects. Our stimuli in a given modality vary in terms of the presented context (2 all, 4 some), target quantifier of the question (3 all 3 some), and the number of objects to which the property in the question applies to in the context (all/none/majority/minority). For instance, if the context is an image with 5 blue blocks and 3 orange blocks, and the question is *Are all blocks in the image blue in color?*, then this would be coded as ‘Some/All/Majority’, since the context shows a scenario where some objects have

the property, the question targets all objects, and the property applies to a majority of the objects in the question. Table 2 shows detailed stimuli.

Vision + Language Stimuli Due to a lack of existing datasets that test knowledge of quantifiers in VLMs in a controlled manner, we design our own Vision and Language dataset, inspired by Gelman et al. (2002) and Smith (1980). To this end, we prompt Gemini 2.5 Flash Image (Comanici et al., 2025) to first generate an image with a specified property that holds true for either all *or* some objects (e.g., flowers in a vase), we then prompt it to edit the image by adding new objects such that the property now applied to the opposite quantifier (e.g., adding a few flowers outside the vase). This gives us our input context variation between *all* and *some*. To ensure the generated images are consistent with our instructions, we manually analyzed *every* generation. Specifically, we prompted the model to generate 200 pairs of images, and then checked if 1) the pairs were minimally different (no visible difference in backgrounds), still maintained

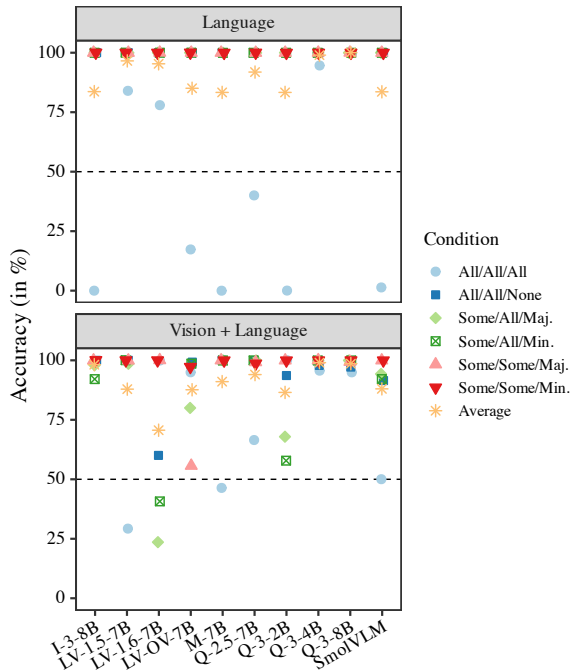


Figure 3: Accuracy of VLMs on *all* and *some*, across modalities. Each point represents the accuracy of a model on a subset of the dataset, except for ‘Average’, which denotes the average accuracy across subsets.

the all vs. some distinction, and 3) did not add additional properties beyond what was specified. This resulted in 140 pairs in total. Appendix C shows full details about our manual annotation. We then paired these images with their respective questions, as described in the previous paragraph (6 different conditions), giving us a total of 840 stimuli. These stimuli mimic those used by Gelman et al. (2002).

Language Stimuli Our language stimuli include synthetically generated contexts posed as word-problems about objects and their properties. We generate 150 unique contexts differing in objects ($N=5$, e.g., *coins in a jar*, *bottles in a crate*, etc.), property pairs ($N=6$, e.g., *red/blue in color*, *contains milk/soda*, etc.), and amounts ($N=5$, e.g., 500 total objects, with 300 having a given property and 200 having the other property). This results in a total of 900 stimuli (when combined with 6 different conditions). The example in Table 2 uses as its object ‘blocks on a table’, as its property pair, ‘red/blue in color’ and as its amount, 20 total objects—12 having the majority property (red) and 8 having the minority property (blue).

4.2 Results and Analysis

Figure 3 shows the accuracy of all models across question-types for each condition, for both modal-

ities. Across both modalities, the VLMs generally achieved high accuracy on questions coded as *Some/Some/**. In contrast, performance on *Some/All/** questions exhibited notable modality-dependent variation. Specifically, all models answered these questions with near-perfect accuracy in the language-only setting, whereas several models were noticeably worse when presented with vision–language inputs, highlighting difficulties that VLMs often have reasoning about “All” and “Some” in the visual domain. More pronounced differences emerged for questions coded as *All/All/All*—a large subset of models performed poorly on these questions, often at or below chance level (especially in the language-only domain), despite maintaining high accuracy on *All/All/None* questions. This pattern suggests a systematic bias toward negative (“no”) responses when models are asked to verify whether a property holds universally across all objects in a scene. Among the evaluated models, only Qwen3 4B and 8B demonstrated near-perfect accuracy across all conditions and modalities.

5 Experiment 3: Constraints on Inductive Inference

Having shown that there exist VLMs (Qwen3-VL 4B and 8B) that can robustly detect categories in images, *as well as* show desirable behavior in distinguishing between *all* and *some*, across modalities, we now turn to our inductive generalization experiments. Here, we test if these VLMs distinguish between propositions that involve explicit quantifiers—*all* and *some*—from those involving *generics*. Following Gelman et al. (2002), we provide models with premises that attribute a property to a category, and varying in the type of proposition (*all/some/generic*) and then ask if a specific member of the category has that property. We then test if VLMs, like humans, show distinct behavior on *all* vs. *generic* vs. *some* (Gelman et al., 2002; Hollander et al., 2002). More specifically, if they show the same qualitative pattern of generalizing the property in the premise to the category member most for *all*, followed by *generic*, followed by *some*. Since the 4B and 8B versions of Qwen3-VL were the only models that *robustly* satisfied our preconditions, we use them for this experiment.

5.1 Data

Category selection We only consider categories for which both models of interest were able to cor-

rectly answer all 20 questions (see §3). Further, most research on generics has largely focused on generic statements about animate concepts—this is presumably because a majority of generics are stated *for* animates (and specifically, animals—see Gelman et al., 2008). To test if the animacy of the premise category has an effect, we sample 50 animate and inanimate categories each ($N_{total}=100$).

Stimuli Our stimuli have two components: a premise and a question. The premise, which follows the same format across our experiments on different modalities, expresses a proposition attributing a property to a category, with the constraints of the scope being modulated by *all*, *some*, or *generic*. To ensure that the influence of the models’ parametric knowledge on their inductive generalization behavior is minimal, we create *novel* properties, using both nonce words (e.g., *is/are daxable*, *has/have feps*, etc.) and more ‘real’ sounding properties (e.g., *has the T9 hormone*, *is made of bergentium*). In total, we have 5 nonce word-based properties and 5 real-sounding properties. Further, we include 4 different prompt variations of expressing the premise (e.g., “*Given that {premise}*”, etc.). This results in a total of 40 different surface form variations—see Appendix D for full details about the surface form variation in our stimuli. The ‘question’ component of the stimuli asks if the property applies to a specific member of the category. For the vision and language stimuli, we do this by using an image followed by a demonstrative—e.g., [image] *Given that all bears have the T9 hormone, does this bear have the T9 hormone?*—following Gelman et al. (2002). For our language stimuli, we use the possessive determiner “*my*”—e.g., *Given that all bears have the T9 hormone, does my bear have the T9 hormone?* All in all, we have 40 surface form variations, 100 categories, and 3 types of premises (all vs. some vs. generic). This gives us 12,000 language stimuli, and 60,000 vision and language stimuli (accounting for 5 images per category).

5.2 Results and Analysis

Main results To evaluate whether VLMs distinguish between *all*, *some*, and *generic* in their inductive generalization, we compute the relative probability of a “yes” response, $p_{rel}(\text{Yes})$, as defined in Eq. (1), for each stimulus. We then average $p_{rel}(\text{Yes})$ across stimuli corresponding to each quantifier category. Figure 4 reports these averages, stratified by modality and animacy, for each model.

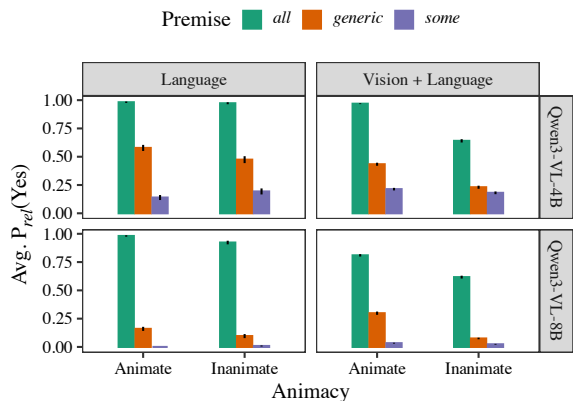


Figure 4: Avg. probabilities of extending the property from the premise to the specific instance in the question (i.e., predicting Yes) in Qwen3-VL models (4B and 8B) for animate and inanimate categories, across both modalities, and premise types (*all* vs. *generic* vs. *some*). Error bars denote 95% confidence intervals, over prompt variations (templates and properties).

In general, models exhibit the highest propensity to attribute a property to an individual category member when the property holds for *all* members of the category, followed by *generic*, followed by *some*. This was also reflected in a linear mixed effects model analysis, predicting $p_{rel}(\text{Yes})$ by using premise type, animacy, and modality as fixed effects, and concept and prompt template as random effects ($p < .001$ for both VLMs). This trend is observed across both modalities, implying that the presence/absence of an image of a member of a category does not affect the models’ qualitative sensitivity to all/some/generic. This indicates a more general representation of *all*, *some*, and *generic*. Overall, the $p_{rel}(\text{Yes})$ in the generics premise is consistently lower for inanimate categories than for animate categories, which aligns with the fact that generics are typically used in the context of animate objects (Gelman et al., 2008). To our knowledge, humans were never tested in a manner similar to that of (Gelman et al., 2002) for inanimate categories, leaving this as a possible direction for future work. Overall, VLMs seem to show qualitatively similar results as humans (Gelman et al., 2002; Hollander et al., 2002), with both systems showing the pattern ALL > GENERICS > SOME.

Post-hoc representation analysis Our behavioral results show that Qwen3 VLMs show distinct patterns of inductive generalization when given premises that express universal, generic, or indefinite propositions; especially for animate kinds. To what extent do these distinctions arise in the inter-




Figure 5: First two Principal Components of the last hidden state representations in selected layers of Qwen3-VL-8B for stimuli that attribute properties to categories and vary in their scope, as modulated by quantifiers (all/every vs. certain/some) or generics (bare plural/indefinite). Results for all layers are shown in Figure 8.

nal representations of the models? Do the Qwen3 VLMs organize propositions in a manner that is consistent with their inductive behavior? To answer these questions, we investigate low-dimensional representations of the models’ hidden states on sentence stimuli that *only* contain the specific proposition, without any task context. By removing task information, we ensure that the patterning of the models’ representation is not biased to mimic their next word-probabilities, which we have already established to be distinguished. A potential confounding factor in this analysis could be that the three premise types are distinguished by surface forms (all vs. some vs. bare plurals), and therefore, differences in the model’s internal representations could simply be explained by surface form differences. In order to combat this, we add a completely different kind of proposition for each premise type. Specifically, we add *every* followed by a singular noun denoting the category to pair with *all*, *certain* to pair with *some*, and an indefinite article ‘a/an’ and a singular noun denoting the category to pair with the bare plural *generic*, as shown in (2):

- (2) a. *Every* bear is daxable. [ALL]
 b. *Certain* bears are daxable. [SOME]
 c. *A bear* is daxable. [GENERIC]

Finally, we restrict our examples to animate categories, and nonce-properties ($N=250$ stimuli per proposition type, and 6 propositions), and extract the models’ hidden state at the final token position across all layers. We then perform principal components analysis (PCA) on these representations. Figure 5 visualizes the first two principal components for layers 1, 6, 12, 18, 24, 30, and 36 (the final layer) of Qwen3-VL 8B.

If the representations are *only* sensitive to surface form as opposed to the similarity in the propositions in terms of their inductive constraints, then we would see six different clusters of points, all

far from each other. Instead, from Figure 5, we see that the representations are largely entangled (in low-dimensional spaces) in the earlier layers, and then gradually organize themselves in terms of their inductive constraints in the middle layers. For instance, in layer 18, we see maximal entanglement in the representations of near-synonymous propositions (*every* with *all*, *certain* with *some*, and indefinite *generic* [*generic-indef*] with bare plural *generic*), as well as maximal separation between propositions with different meanings. This in fact starts to emerge at layer 13 (see Figure 8 in the appendix). In later layers, while the separation among propositions with different meanings is further reduced, propositions similar in meaning are still closer to each other than they are to other propositions. We provide quantitative analysis of this effect in Appendix E. Overall, these results suggest that the propositions differing in their inductive constraints—as modulated by quantifiers or generics—are represented separately in VLMs’ representational space, in a manner that cannot be explained by surface-form differences alone.

6 Conclusion

The *form* of language can have systematic impact on its *function*. We have presented an analysis of how this could be realized within VLMs, by focusing on how subtle linguistic cues that modulate the scope of a proposition affect these models’ inductive inferences involving categories. Our findings suggest that VLMs represent propositions varying in these cues in ways that are similar to those in humans (Gelman et al., 2002). That is, there was general behavioral alignment between how *all*, *some*, and *generics* constrain the inductive behavior between both systems (humans and models). Overall, since generics in particular have been a great mystery to philosophers, cognitive scientists, and linguists (Gelman, 2004; Leslie, 2008),

we are hopeful that further analyses on (V)LMs can help shed light on the underlying mechanisms for the acquisition and processing of generic statements (Allaway et al., 2024).

Limitations

This study is a foray into understanding the connection between generics and inductive inferences from a computational perspective. While we were able to establish behavioral alignment between VLMs and humans, our study has several limitations, which we enumerate below.

Lack of Causal Implications First, while we have shown generics to be represented separately from quantifiers like *all* and *some* in a manner that is not just determined by surface form differences, we have not established any causal implications of such representations. In general there have been cautions against exclusively concluding from simpler low-dimensional analyses, since they may be biased towards certain kinds of features (Lampinen et al., 2025). A more thorough investigation using causal interpretability methods (Geiger et al., 2024) could potentially alleviate this issue, though it is not entirely clear how one might causally localize regions of interest pertaining to the three types of propositions in a manner that does not encode differences between them—ideally we would want to discover this in an unsupervised manner.

Model Choice and ‘Cognitive Plausibility’ Perhaps one of the biggest limitations of this work is that the VLMs we investigate have been trained on orders of magnitude more data than humans, which raises questions about the transfer of insights between models and humans. While we have refrained from making any claims about human cognition from our analyses, especially due to this fact, we have enforced strict preconditions in order to select the right model to investigate and have remained faithful to the pretests of Gelman et al. (2002), which this paper aims to replicate. At the same time, a recent working hypothesis, called the *contravariance principle*, has emerged in the sub-community of researchers interested in establishing a transfer of insight between neural networks and human minds (Cao and Yamins, 2024a,b; Futrell and Mahowald, 2025). According to this principle, if two systems seemingly solve the same problem that is *hard*—the sense that solving it requires satis-

fyng innumerable constraints—then their solutions are similar in “functionally relevant” ways, even though the systems themselves might differ from one another in “functionally irrelevant ways” (Cao and Yamins, 2024b). Since generics in particular has been one of the hardest puzzles in the cognitive psychology of language and thought (Leslie, 2008), there is some benefit to first establish that the two systems that we have concerned ourselves with (VLMs and Humans) converge at the same behavioral sensitivities. In this way, we have established which models are better animal models that we can then investigate more mechanistically in order to articulate novel predictions and hypotheses (McCloskey, 1991; Lakretz et al., 2021; Misra and Kim, 2024).

Sample Sizes for All and Some Our benchmark for establishing behavioral sensitivity to *all* and *some*—inspired from developmental work though it may be—only has a few hundred instances, limited to a handful of attributes and objects. While the benchmark is not the main contribution of the paper and serves more as a pre-test for models, it would be worthwhile to increasing the benchmark’s sample size and diversity, as well as include an entire host of quantifiers, in future work.

Model generated stimuli for All and Some Our *all* and *some* pretest involves both language and vision+language stimuli. While we generated our language stimuli ourselves, we had to rely on Gemini 2.5 (Comanici et al., 2025) to generate image stimuli for our vision+language subset. The main reason behind this was that there were no controlled datasets that allowed us to test models’ behavior on *all* and *some* in the same manner as Smith (1980) and Gelman et al. (2002). Furthermore, it is expensive—intractable, even—to find images in the world that could serve as controlled stimuli that differ only minimally in the applicability of a quantifier. Finally, we held out stimuli towards the strictest of standards, and manually inspected each and every single generated image. We ended up rejecting about 30% of the total generated image-pairs, and have disclosed full details in §C. Our usage of Gemini to generate stimuli for this benchmark is another reason why our sample sizes are small (see above limitation).

Monolingual Investigation Finally, our stimuli only focus on the English language, whereas generics have been a universal puzzle across languages,

because they are never explicitly marked (Leslie, 2008). Understanding them from a multi-lingual perspective would make for important future work.

Acknowledgments

We are grateful to Susan Gelman and her group for their work on generics, kinds, and inductive inferences, which is the important foundational work we build on. We are also grateful to the THINGS initiative (Hebart et al., 2019, 2023) for making their data open and accessible. K.M. is supported by the Donald D. Harrington Faculty Fellowship at UT Austin. The authors acknowledge the Texas Advanced Computing Center (TACC) at The University of Texas at Austin for providing computational resources that have contributed to the research results reported within this paper. URL: <http://www.tacc.utexas.edu>.

References

- Emily Allaway, Chandra Bhagavatula, Jena D Hwang, Kathleen McKeown, and Sarah-Jane Leslie. 2024. Exceptions, instantiations, and overgeneralization: Insights into how language models process generics. *Computational Linguistics*, 50(4):1211–1275.
- Emily Allaway, Jena D. Hwang, Chandra Bhagavatula, Kathleen McKeown, Doug Downey, and Yejin Choi. 2023. [Penguins don't fly: Reasoning about generics through instantiations and exceptions](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2618–2635, Dubrovnik, Croatia. Association for Computational Linguistics.
- Emily M. Bender and Alexander Koller. 2020. [Climbing towards NLU: On meaning, form, and understanding in the age of data](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5185–5198, Online. Association for Computational Linguistics.
- Sudeep Bhatia. 2023. Inductive reasoning in minds and machines. *Psychological Review*.
- Rosa Cao and Daniel Yamins. 2024a. Explanatory models in neuroscience, part 1: Taking mechanistic abstraction seriously. *Cognitive Systems Research*, 87:101244.
- Rosa Cao and Daniel Yamins. 2024b. Explanatory models in neuroscience, part 2: Functional intelligibility and the contravariance principle. *Cognitive Systems Research*, 85:101200.
- Gustavo Cilleruelo, Emily Allaway, Barry Haddow, and Alexandra Birch. 2025. [Generics are puzzling. can language models find the missing piece?](#) In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 6571–6588, Abu Dhabi, UAE. Association for Computational Linguistics.
- Andrei Cimpian and Ellen M Markman. 2008. Preschool children's use of cues to generic meaning. *Cognition*, 107(1):19–53.
- Claudia Collacciani, Giulia Rambelli, and Marianna Bolognesi. 2024. [Quantifying generalizations: Exploring the divide between human and LLMs' sensitivity to quantification](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11811–11822, Bangkok, Thailand. Association for Computational Linguistics.
- Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, and 1 others. 2025. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*.
- Matt Deitke, Christopher Clark, Sangho Lee, Rohun Tripathi, Yue Yang, Jae Sung Park, Mohammadreza Salehi, Niklas Muennighoff, Kyle Lo, Luca Soldaini, and 1 others. 2025. Molmo and pixmo: Open weights and open data for state-of-the-art vision-language models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 91–104.
- Chaz Firestone. 2020. Performance vs. competence in human-machine comparisons. *Proceedings of the National Academy of Sciences*, 117(43):26562–26571.
- Stella Frank and Emily Allaway. 2025. [VISaGE: Understanding visual generics and exceptions](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 32537–32546, Suzhou, China. Association for Computational Linguistics.
- Richard Futrell and Kyle Mahowald. 2025. How linguistics learned to stop worrying and love the language models. *arXiv preprint arXiv:2501.17047*.
- Atticus Geiger, Zhengxuan Wu, Christopher Potts, Thomas Icard, and Noah Goodman. 2024. Finding alignments between interpretable causal variables and distributed neural representations. In *Causal Learning and Reasoning*, pages 160–187. PMLR.
- Susan A Gelman. 2004. Learning words for kinds: Generic noun phrases in acquisition. *Weaving a lexicon*, pages 445–484.
- Susan A Gelman, Peggy J Goetz, Barbara W Sarnecka, and Jonathan Flukes. 2008. Generic language in parent-child conversations. *Language Learning and Development*, 4(1):1–31.

- Susan A Gelman, Jon R Star, and Jonathan Flukes. 2002. Children’s use of generics in inductive inferences. *Journal of Cognition and Development*, 3(2):179–199.
- Simon Jerome Han, Keith J Ransom, Andrew Perfors, and Charles Kemp. 2024. Inductive reasoning in humans and large language models. *Cognitive Systems Research*, 83:101155.
- Martin N Hebart, Oliver Contier, Lina Teichmann, Adam H Rockter, Charles Y Zheng, Alexis Kidder, Anna Corriveau, Maryam Vaziri-Pashkam, and Chris I Baker. 2023. [Things-data, a multimodal collection of large-scale datasets for investigating object representations in human brain and behavior](#). *eLife*, 12:e82580.
- Martin N Hebart, Adam H Dickter, Alexis Kidder, Wan Y Kwok, Anna Corriveau, Caitlin Van Wicklin, and Chris I Baker. 2019. Things: A database of 1,854 object concepts and more than 26,000 naturalistic object images. *PLoS one*, 14(10):e0223792.
- Michelle A Hollander, Susan A Gelman, and Jon Star. 2002. Children’s interpretation of generic noun phrases. *Developmental psychology*, 38(6):883.
- Philipp Kaniuth, Florian P Mahner, Jonas Perkuhn, and Martin N Hebart. 2024. A high-throughput approach for the efficient prediction of perceived similarity of natural objects. *bioRxiv*, pages 2024–06.
- Yair Lakretz, Dieuwke Hupkes, Alessandra Vergallito, Marco Marelli, Marco Baroni, and Stanislas Dehaene. 2021. Mechanisms for handling nested dependencies in neural-network language models and humans. *Cognition*, 213:104699.
- Andrew Kyle Lampinen, Stephanie CY Chan, Yuxuan Li, and Katherine Hermann. 2025. Representation biases: will we achieve complete understanding by analyzing representations? *arXiv preprint arXiv:2507.22216*.
- Hugo Laurençon, Andrés Marafioti, Victor Sanh, and Léo Tronchon. 2024. [Building and better understanding vision-language models: insights and future directions](#). *Preprint*, arXiv:2408.12637.
- Sarah-Jane Leslie. 2008. Generics: Cognition and acquisition. *Philosophical review*, 117(1):1–47.
- Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Peiyuan Zhang, Yanwei Li, Ziwei Liu, and Chunyuan Li. 2025. [LLaVA-onevision: Easy visual task transfer](#). *Transactions on Machine Learning Research*.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2023a. Improved baselines with visual instruction tuning.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023b. Visual instruction tuning. In *NeurIPS*.
- Kyle Mahowald, Anna A Ivanova, Idan A Blank, Nancy Kanwisher, Joshua B Tenenbaum, and Evelina Fedorenko. 2024. Dissociating language and thought in large language models. *Trends in cognitive sciences*.
- Andrés Marafioti, Orr Zohar, Miquel Farré, Merve Noyan, Elie Bakouch, Pedro Cuenca, Cyril Zakka, Loubna Ben Allal, Anton Lozhkov, Nouamane Tazi, Vaibhav Srivastav, Joshua Lochner, Hugo Larcher, Mathieu Morlon, Lewis Tunstall, Leandro von Werra, and Thomas Wolf. 2025. Smolvlm: Redefining small and efficient multimodal models. *arXiv preprint arXiv:2504.05299*.
- Michael McCloskey. 1991. Networks and theories: The place of connectionism in cognitive science. *Psychological science*, 2(6):387–395.
- Kanishka Misra. 2022. minicons: Enabling flexible behavioral and representational analyses of transformer language models. *arXiv:2203.13112*.
- Kanishka Misra, Allyson Ettinger, and Julia Rayz. 2021. Do language models learn typicality judgments from text? In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 43.
- Kanishka Misra and Najoung Kim. 2024. Generating novel experimental hypotheses from language models: A case study on cross-dative generalization. *arXiv preprint arXiv:2408.05086*.
- Kanishka Misra, Julia Rayz, and Allyson Ettinger. 2022. A property induction framework for neural language models. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 44.
- Kanishka Misra, Julia Rayz, and Allyson Ettinger. 2023. [COMPS: Conceptual minimal pair sentences for testing robust property knowledge and its inheritance in pre-trained language models](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2928–2949, Dubrovnik, Croatia. Association for Computational Linguistics.
- Gregory Murphy. 2004. *The big book of concepts*. MIT press.
- Daniel N Osherson, Edward E Smith, Ormond Wilkie, Alejandro Lopez, and Eldar Shafir. 1990. Category-based induction. *Psychological review*, 97(2):185.
- Sandeep Prasada. 2000. Acquiring generic knowledge. *Trends in cognitive sciences*, 4(2):66–72.
- Yulu Qin, Dheeraj Varghese, Adam Dahlgren Lindström, Lucia Donatelli, Kanishka Misra, and Najoung Kim. 2025. [Vision-and-language training helps deploy taxonomic knowledge but does not fundamentally alter it](#). In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Qwen Team. 2025a. [Qwen2.5-vl](#).
- Qwen Team. 2025b. [Qwen3 technical report](#). *Preprint*, arXiv:2505.09388.

Sello Ralethe and Jan Buys. 2022. [Generic overgeneralization in pre-trained language models](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 3187–3196, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.

Juan Diego Rodriguez, Aaron Mueller, and Kanishka Misra. 2025. [Characterizing the role of similarity in the property inferences of language models](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 11515–11533, Albuquerque, New Mexico. Association for Computational Linguistics.

Carol L Smith. 1980. Quantifiers and question answering in young children. *Journal of experimental child psychology*, 30(2):191–205.

Alvin Wei Ming Tan, Chunhua Yu, Bria Lorelle Long, Wanjing Anya Ma, Tonya Murray, Rebecca D. Silverman, Jason D Yeatman, and Michael Frank. 2024. [Devbench: A multimodal developmental benchmark for language learning](#). In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.

Hugh Mee Wong, Rick Nouwen, and Albert Gatt. 2025. [Vaquum: Are vague quantifiers grounded in visual data?](#) *arXiv preprint arXiv:2502.11874*.

Eunice Yiu, Maan Qraitem, Anisa Noor Majhi, Charlie Wong, Yutong Bai, Shiry Ginosar, Alison Gopnik, and Kate Saenko. 2025. [KiVA: Kid-inspired visual analogies for testing large multimodal models](#). In *The Thirteenth International Conference on Learning Representations*.

Charles Y. Zheng, Francisco Pereira, Chris I. Baker, and Martin N. Hebart. 2019. [Revealing interpretable object representations from human behavior](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.

A Selected Models

Table 3 shows the metadata of the models used in this paper. Models were run on a combination of NVIDIA A40 and GH100 GPUs.

B Data Release

We used images from the THINGS database for our vision+language stimuli (Hebart et al., 2019, 2023). They release their data with the Attribution CC BY license.² When we release our stimuli we will point to their file names and not release their data as part of our data. All our data will be released under the MIT License.

²See <https://osf.io/jum2f/files/52wrx>

C Vision + Language Stimuli Generation for all and some

C.1 Prompts

We prompt Gemini 2.5 Flash Image (Comanici et al., 2025) to first generate an image, and then prompt the model to modify the generated image. All in all, this cost around \$17 in google cloud credits. In the first condition, we first ask the model to generate an image with all the objects in the image are in the same color, and then modify the image by adding some objects in a different color:

1:

USER:
Create a picture of {number_majority} {color} {object}. Make sure the background is as clean as possible.

MODEL: [IMAGE_ALL]

2:

USER:
[IMAGE_ALL] Modify this picture by adding {number_minority} {color} {object}. Change nothing else.

MODEL:
[IMAGE_SOME]

In the second condition, we first prompt the model to generate an image with all the objects in the image in a container, and then prompt it to modify the image by adding some objects (minority) out of the container:

1:

USER:
Generate a picture of {number_majority} {object} in a {container}. Make sure the background is as clean as possible.

MODEL: [IMAGE_ALL]

2:

USER:
[IMAGE_ALL] Modify this picture by adding {number_minority} {object} outside the {container}. Change nothing else.

MODEL:
[IMAGE_SOME]

In the third condition, we first prompt the model to generate an image with most of the objects (majority) in a container, and some of them (minority) out of it. Then we prompt the model to modify the

Model	Params	Instruction Tuned	HF Identifier
SmolVLM	2.2B	✓	SmolVLM-Instruct
LLaVA-1.5	7B	✓	llava-1.5-7b-hf
LLaVA-1.6 (Mistral)	7B	✓	llava-v1.6-mistral-7b-hf
LLaVA-OneVision (Qwen)	7B	✓	llava-hf/llava-onevision-qwen2-7b-ov-hf
Molmo	7B	✓	Molmo-7B-D-0924
Qwen2.5-VL	7B	✓	Qwen2.5-VL-7B-Instruct
Qwen3-VL	2B	✓	Qwen3-VL-2B-Instruct
Qwen3-VL	4B	✓	Qwen3-VL-4B-Instruct
Qwen3-VL	8B	✓	Qwen3-VL-8B-Instruct

Table 3: Overview of models used in our experiments.

image by removing the objects that are out of the container.

1:

USER:

Generate a picture of {number_majority} {object} in a {container} and {number_minority} out of it. Make sure the background is as clean as possible.

MODEL: [IMAGE_SOME]

2:

USER:

[IMAGE_SOME] Modify this picture by removing the {number_minority} {object} outside the {container}. Change nothing else.

MODEL:

[IMAGE_ALL]

C.2 Annotation

We manually annotate the 200 generated image pairs using the frontend shown in Figure 6. We excluded 60 pairs from our dataset. Examples of excluded pairs are shown in Figure 7.

D Surface form variation in the inductive generalization stimuli

In order to robustly identify how tested VLMs differentiate between *generics*, *all* and *some*, we use 4 different surface form variations for each stimuli type-category combination. Table 4 shows the different surface form variations used for the stimuli in Experiment 3 along with examples of each.

E Detailed PCA results

Figure 8 shows low dimensional representations of the stimuli in our post-hoc analysis in §5.2. We see strong evidence of inductive-behavior-based separation of points starting to emerge at layer 13.

To further quantify the degree of closeness between classes of stimuli, we compute the average pairwise euclidean distances of each point to every other point, and visualize results by taking the average euclidean distance between the set of points within each type of premise. We see that the average distance between points belonging to propositions with similar inductive constraints (*generic-generic-indef*, *all-every*, *some-certain*) is lowest, relative to all others, and this pattern emerges around layer 13, and becomes stable at around layer 22.

Pair Verification

Pair 1 of 200 · 199 judged




Image 0




Image 1

```

scenario colors_all_some
noun birds
colors [
  "red",
  "blue"
]
counts [
  6,
  3
]
prompts [
  "Create a picture of
6 red birds. Make sure
the background is as
clean as possible.",
  "Modify this picture
by adding 3 blue birds.
Change nothing else."
]
model gemini-2.5-flash-image-
preview
created_at 2025-10-
14T20:20:30.877343
responses [
  "Here's your picture
of six red birds with a
clean background: ",
  "Here's your modified
picture with 3 blue
        
```

Previous
Next
Jump to pair: 1 - pair_0000_10-14
Mark as Valid
Mark as Invalid
Clear
Download Judgments

Figure 6: Frontend for annotating image pairs generated by Gemini-2.5

Template Name	Prompt
Given that	Given that all bears have the T9 Hormone, does this bear have the T9 Hormone?
Suppose	Suppose all bears have the T9 Hormone. Does this bear have the T9 Hormone?
If	If all bears have the T9 Hormone, does this bear have the T9 Hormone?
No Template	All bears have the T9 Hormone. Does this bear have the T9 Hormone?

Table 4: Surface Form Variations for Experiment 3

Type	Animacy	Property
Real-sounding	Animate	has the T9 Hormone / has Interleukin 54 / has Homoglobin 76 / has Benzokinase / has Immunoglobulin P
	Inanimate	is made of bergentium, is made of misrium, is heteromagnetic, is betamagnetic, contains ostinium
Nonce	Animate/Inanimate	is/are daxable, has/have feps, can dax, can gek, is/are totable

Table 5: Properties for Experiment 3. The nonce properties are used universally for Animate and Inanimate concepts.



(1a)



(1b)



(2a)



(2b)

Figure 7: Examples of invalid image pairs. For the first pair, the model was prompted to generate an image with six coins in the bag and three coins outside the bag in (1a), and then to remove the three coins outside the bag in (1b); however, the model did not follow the prompt, and 2 coins still appear outside the bag in (1b). For the second pair, the model was prompted to generate an image with 6 orange balls in (2a), and then to add three brown balls in (2b); however, the model changed the color of balls when prompted to modify the image.

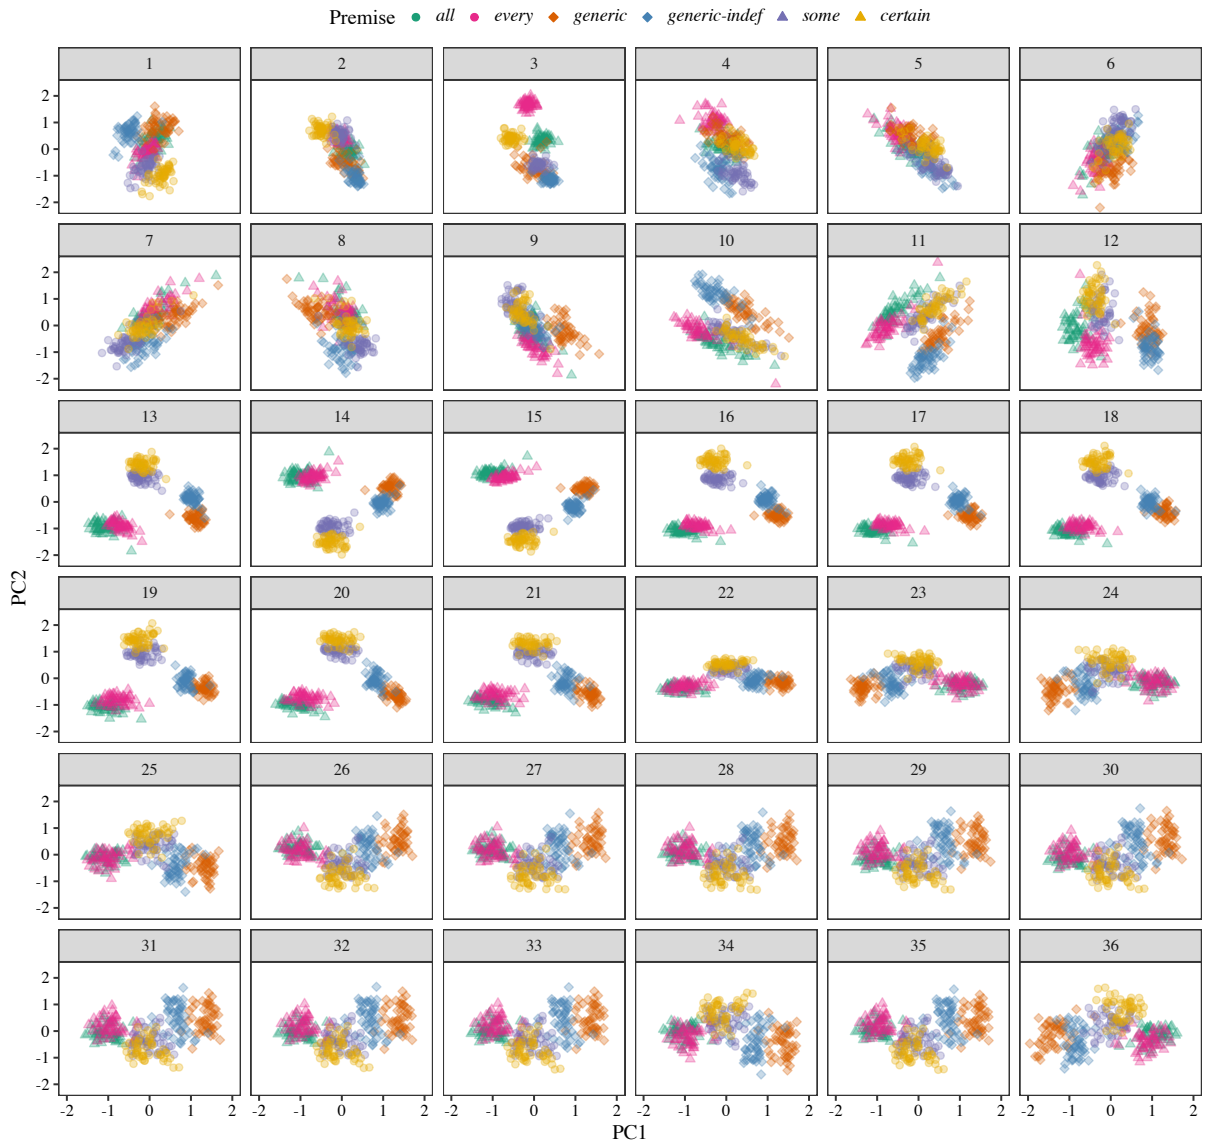


Figure 8: First and second Principal Components of the last hidden state representations in all layers of Qwen3-VL-8B for stimuli that attribute properties to categories and vary in their scope, as modulated by quantifiers (all/every vs. some/certain) or generics (bare plural/indefinite).




Figure 9: Average euclidean distance between collection of points reduced to two dimensions (using PCA). The average distance between points belonging to propositions with similar inductive constraints (*generic-generic-indef*, *all-every*, *some-certain*) is lowest, relative to all others, and those emerges around layer 13.